



Google Research Blog

The latest news from Research at Google

Sudoku, Linear Optimization, and the Ten Cent Diet

Posted: Tuesday, September 30, 2014



807

Tweet

96

Like

175

Posted by Jon
Orwant,
Engineering Manager

(cross-posted on the [Google Apps Developer blog](#), and the [Google Developers blog](#))

In 1945, future Nobel laureate [George Stigler](#) wrote an essay in the Journal of Farm Economics titled *The Cost of Subsistence* about a seemingly simple problem: how could a soldier be fed for as little money as possible?

The “Stigler Diet” became a classic problem in the then-new field of [linear optimization](#), which is used today in many areas of science and engineering. Any time you have a set of linear constraints such as “at least 50 square meters of solar panels” or “the amount of paint should equal the amount of primer” along with a linear goal (e.g., “minimize cost” or “maximize customers served”), that’s a linear optimization problem.

At Google, our engineers work on plenty of optimization problems. One example is our [YouTube video stabilization system](#), which uses linear optimization to eliminate the shakiness of handheld cameras. A more lighthearted example is in the [Google Docs Sudoku add-on](#), which instantaneously generates and solves Sudoku puzzles inside a Google Sheet, using the [SCIP](#) mixed integer programming solver to compute the solution.

Research at Google

google.com/+ResearchatGoogle

▼ x, CS+x



Follow

+1

+ 867,967

[Labels](#)

[Archive](#)

[Feed](#)

[Follow @googleresearch](#)

Give us feedback in our [Product Forums](#).



Add-ons

Check out [the store](#) for many more

Sudoku Sheets
★★★★★ (18) · [from sudokusheets](#) · 4,443 users

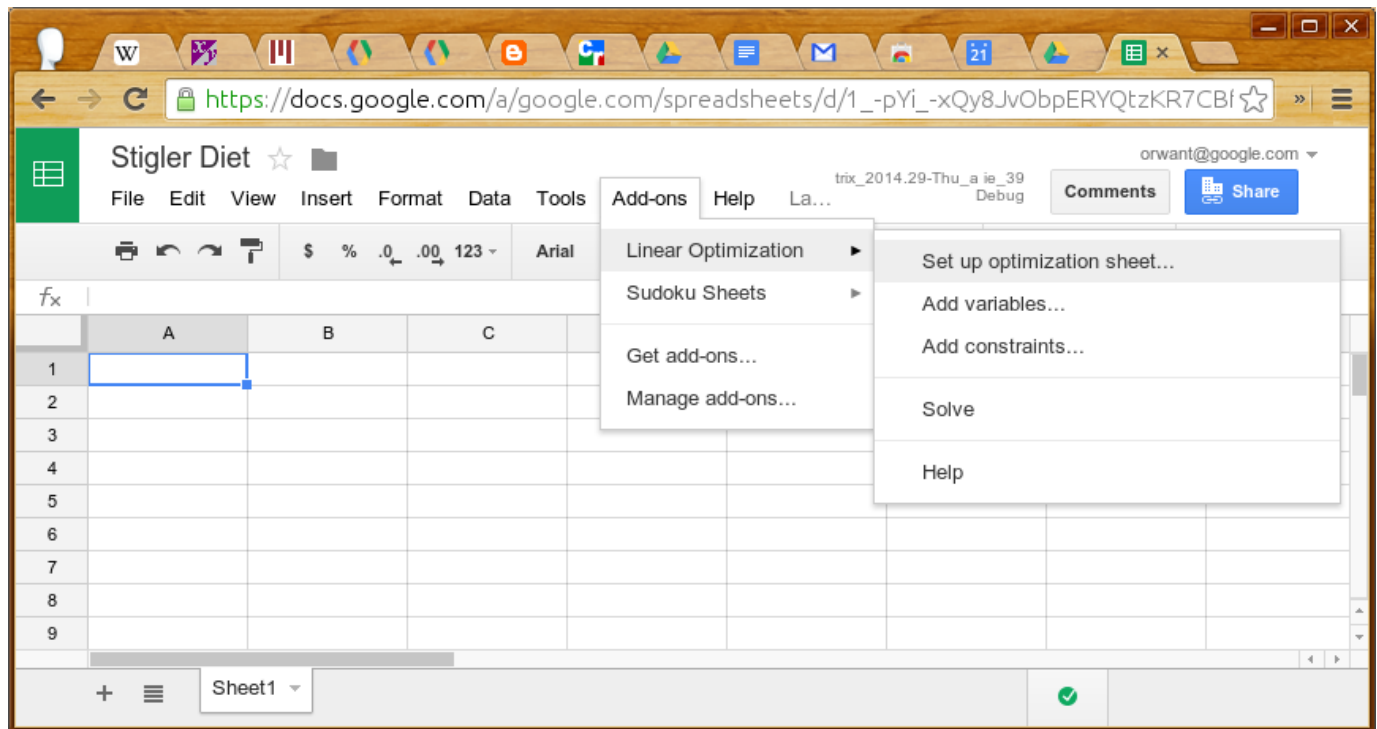
INSTALLED

Sudoku Sheets menu:

- Generate sudoku
- Disable Assisted Mode
- Check puzzle
- Solve
- Clear solution
- Help

	A	B	C	D	E	F	G	H	I
1	Input Grid	3	6			9			2
2							8		
3			9		5				1
4								5	
5					1	6	3	7	9
6				9					6
7									4
8				7		2		1	8
9			5	1		7			9
10									
11	Solution Grid	3	6	8	5	9	1	4	2
12		7	2	1	6	4	8	9	3
13		9	4	5	2	3	7	6	8
14		6	3	4	8	7	9	5	1
15		8	5	2	1	6	3	7	4
16		1	7	9	4	5	2	8	6

Today we're proud to announce two new ways for everyone to solve linear optimization problems. First, you can now solve linear optimization problems in Google Sheets with the [Linear Optimization add-on](#) written by Google Software Engineer Mihai Amarandei-Stavila. The add-on uses Google Apps Script to send optimization problems to Google servers. The solutions are displayed inside the spreadsheet. For developers who want to create their own applications on top of Google Apps, we also provide an [API](#) to let you call our linear solver directly.



Second, we're open-sourcing the linear solver underlying the add-on: Glop (the Google Linear Optimization Package), created by [Bruno de Backer](#) with other members of the Google Optimization team. It's available as part of the [or-tools suite](#) and we provide a [few examples](#) to get you started. On that page, you'll find the Glop solution to the Stigler diet problem. (A Google Sheets file that uses Glop and the Linear Optimization add-on to solve the Stigler diet problem is available [here](#). You'll need to [install the add-on first](#).)

Stigler posed his problem as follows: given nine nutrients (calories, protein, Vitamin C, and so on) and 77 candidate foods, find the foods that could sustain soldiers at minimum cost.

The [Simplex algorithm](#) for linear optimization was two years away from being invented, so Stigler had to do his best, arriving at a diet that cost \$39.93 per year (in 1939 dollars), or just over ten cents per day. Even that wasn't the cheapest diet. In 1947, Jack Laderman used Simplex, nine calculator-wielding clerks, and 120 person-days to arrive at the optimal solution.

Glop's Simplex implementation solves the problem in 300 milliseconds. Unfortunately, Stigler didn't include taste as a constraint, and so the poor hypothetical soldiers will eat nothing but the following, ever:

- Enriched wheat flour
- Liver

- Cabbage
- Spinach
- Navy beans

Is it possible to create an appealing dish out of these five ingredients? Google Chef Anthony Marco took it as a challenge, and we're calling the result *Foie Linéaire à la Stigler*.



This optimal meal consists of seared calf liver dredged in flour, atop a navy bean purée with marinated cabbage and a spinach pesto.

Chef Marco reported that the most difficult constraint was making the dish tasty without butter or cream. That said, I had the opportunity to taste our linear optimization solution, and it was delicious.

Labels: [Linear Optimization](#), [open source](#)
[110 comments](#)

Collaborative Mathematics with SageMathCloud and Google Cloud

Platform

Posted: Monday, September 29, 2014

 412

 71

 75

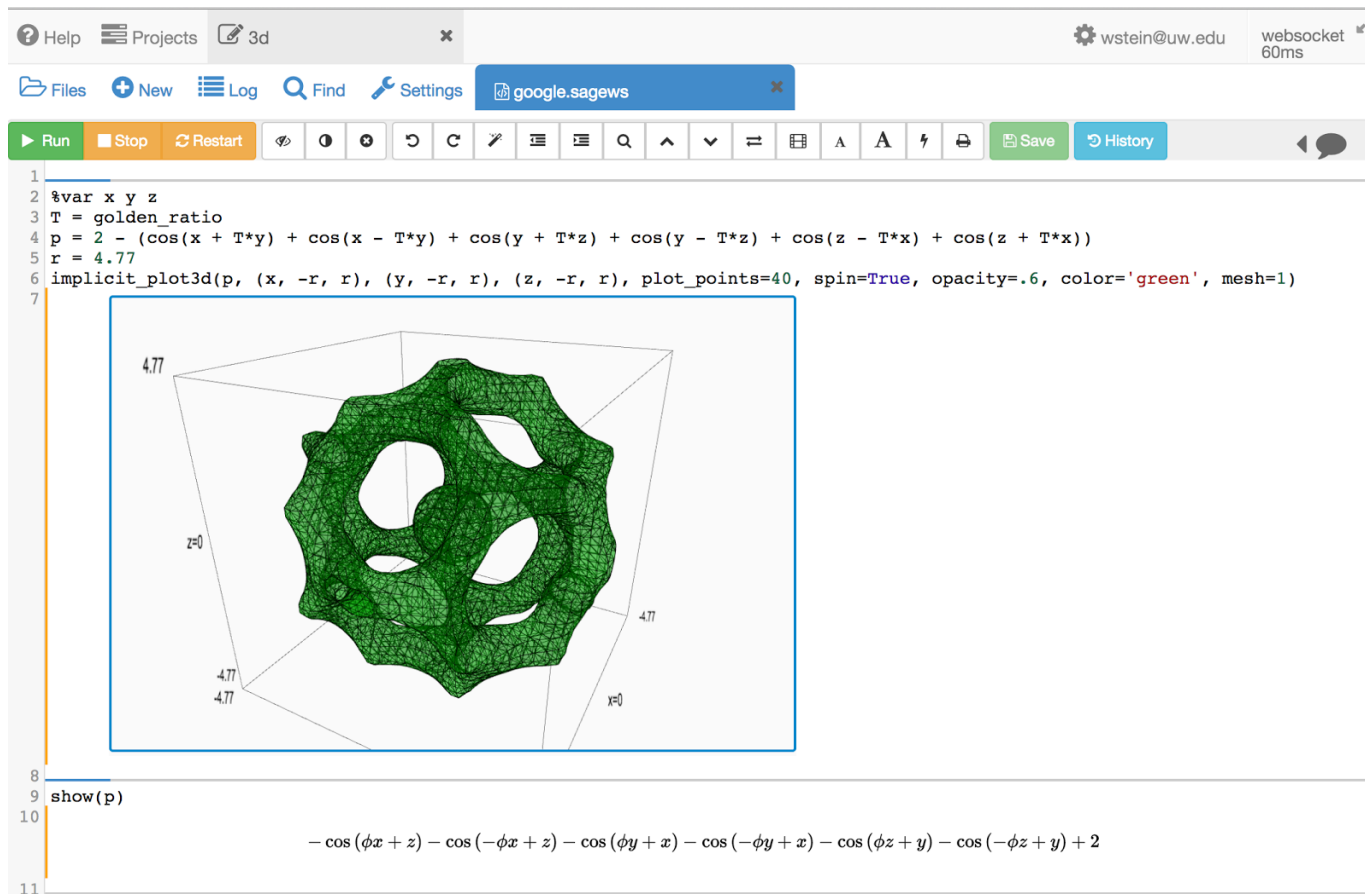
Posted by
Craig Citro,
Software Engineer

(cross-posted on the [Google for Education blog](#) and [Google Cloud Platform blog](#))

Modern mathematics research is distinguished by its openness. The notion of "mathematical truth" depends on theorems being published with proof, letting the reader understand how new results build on the old, all the way down to basic mathematical axioms and definitions. These new results become tools to aid further progress.

Nowadays, many of these tools come either in the form of software or theorems whose proofs are supported by software. If new tools produce unexpected results, researchers must be able to collaborate and investigate how those results came about. Trusting software tools means being able to inspect and modify their source code. Moreover, open source tools can be modified and extended when research veers in new directions.

In an attempt to create an open source tool to satisfy these requirements, University of Washington Professor [William Stein](#) built [SageMathCloud](#) (or SMC). SMC is a robust, low-latency web application for collaboratively editing mathematical documents and code. This makes SMC a viable platform for mathematics research, as well as a powerful tool for teaching any mathematically-oriented course. SMC is built on top of standard open-source tools, including [Python](#), [LaTeX](#), and [R](#). In 2013, William received a 2013 Google Research Award which provided [Google Cloud Platform](#) credits for SMC development. This allowed William to extend SMC to use [Google Compute Engine](#) as a hosting platform, achieving better scalability and global availability.



SMC allows users to interactively explore 3D graphics with only a browser

SMC has its roots in 2005, when William started the [Sage](#) project in an attempt to create a viable free and open source alternative to existing closed-source mathematical software. Rather than starting from scratch, Sage was built by making the best existing open-source mathematical software work together transparently and filling in any gaps in functionality.

During the first few years, Sage grew to have about 75K active users, while the developer community matured with well over 100 contributors to each new Sage release and about 500 developers contributing [peer-reviewed code](#).

Inspired by Google Docs, William and his students built the first web-based interface to Sage in 2006, called [The Sage Notebook](#). However, The Sage Notebook was designed for a small number of users and would work for a small group (such as a single class), but soon became difficult to maintain for larger groups, let alone the whole web.

As the growth of new users for Sage began to stall in 2010, due largely to installation complexity, William turned his attention to finding ways to expand Sage's availability to a broader audience. Based on his experience teaching his own courses with Sage,

and feedback from others doing the same, William began building a new Web-hosted version of Sage that can scale to the next generation of users.

The result is [SageMathCloud](#), a highly distributed multi-datacenter application that creates a viable way to do computational mathematics collaboratively online. SMC uses a wide variety of open source tools, from languages ([CoffeeScript](#), [node.js](#), and [Python](#)) to infrastructure-level components (especially [Cassandra](#), [ZFS](#), and [bup](#)) and a number of in-browser toolkits (such as [CodeMirror](#) and [three.js](#)).

Latency is critical for collaborative tools: like an online video game, everything in SMC is interactive. The initial versions of SMC were hosted at UW, at which point the distance between Seattle and far away continents was a significant issue, even for the fastest networks. The global coverage of Google Cloud Platform provides a low-latency connection to SMC users around the world that is both fast and stable. It's not uncommon for long-running research computations to last days, or even weeks -- and here the robustness of Google Compute Engine, with machines live-migrating during maintenance, is crucial. Without it, researchers would often face multiple restarts and delays, or would invest in engineering around the problem, taking time away from the core research.

SMC sees use across a number of areas, especially:

- **Teaching:** any course with a programming or math software component, where you want all your students to be able to use that component without dealing with the installation pain. Also, SMC allows students to easily share files, and even work together in realtime. There are [dozens of courses](#) using SMC right now.
- **Collaborative Research:** all co-authors of a paper can work together in an SMC project, both writing the paper there and doing research-level computations.

Since launching SMC in May 2013, there are already more than 20,000 monthly active users who've started using Sage via SMC. We look forward to seeing if SMC has an impact on the number of active users of Sage, and are excited to learn about the collaborative research and teaching that it makes possible.

Labels: [Cloud Computing](#), [open source](#), [Research Awards](#)
[77 comments](#)

Introducing Structured Snippets, now a part of Google Web Search

Posted: Monday, September 22, 2014

 774

 299

 346

Posted by

Corinna Cortes,

Boulos Harb, Afshin Rostamizadeh, Ken Wilder, and Cong Yu,

Google Research

Google Web Search has evolved in recent years with a host of features powered by the [Knowledge Graph](#) and other data sources to provide users with highly structured and relevant data. Structured Snippets is a new feature that incorporates facts into individual result snippets in Web Search. As seen in the example below, interesting and relevant information is extracted from a page and displayed as part of the snippet for the query “*nikon d7100*”:

Nikon D7100 In-Depth Review: Digital Photography Review

www.dpreview.com/reviews/nikon-d7100 ▾ DP Review ▾

Review based on a production **Nikon D7100** running firmware 1.0. The enthusiast-targeted **Nikon D7100** becomes the company's latest APS-C DSLR to feature ...

Sensor resolution (type): 24MP CMOS ... **Weight (with battery):** 765 g (1 lb, 11 oz)
Display size / resolution: Fixed 3.2", 1.2...

The WebTables research team has been working to extract and understand tabular data on the Web with the intent to surface particularly relevant data to users. Our data is already used in the [Research Tool found in Google Docs and Slides](#); Structured Snippets is the latest collaboration between Google Research and the Web Search team employing that data to seamlessly provide the most relevant information to the user. We use machine learning techniques to distinguish data tables on the Web from uninteresting tables, e.g., tables used for formatting web pages. We also have additional algorithms to determine quality and relevance that we use to display up to four highly ranked facts from those data tables. Another example of a structured snippet for the query “*superman*”, this time as it appears on a mobile phone, is shown below:

Superman - Wikipedia, the free encyclopedia

en.m.wikipedia.org/wiki/Superman

Superman is the superhero persona of a fictional character that appears in comic books published by ...

Created by: Jerry Siegel; Joe Shuster

Place of origin: Krypton

First appearance: Action Comics #1; ...

Fact quality will vary across results based on page content, and we are continually enhancing the relevance and accuracy of the facts we identify and display. We hope users will find this extra snippet information useful.




Labels: [Search](#), [Structured Data](#)

[203 comments](#)

Sign in to edx.org with Google (and Facebook, and...)

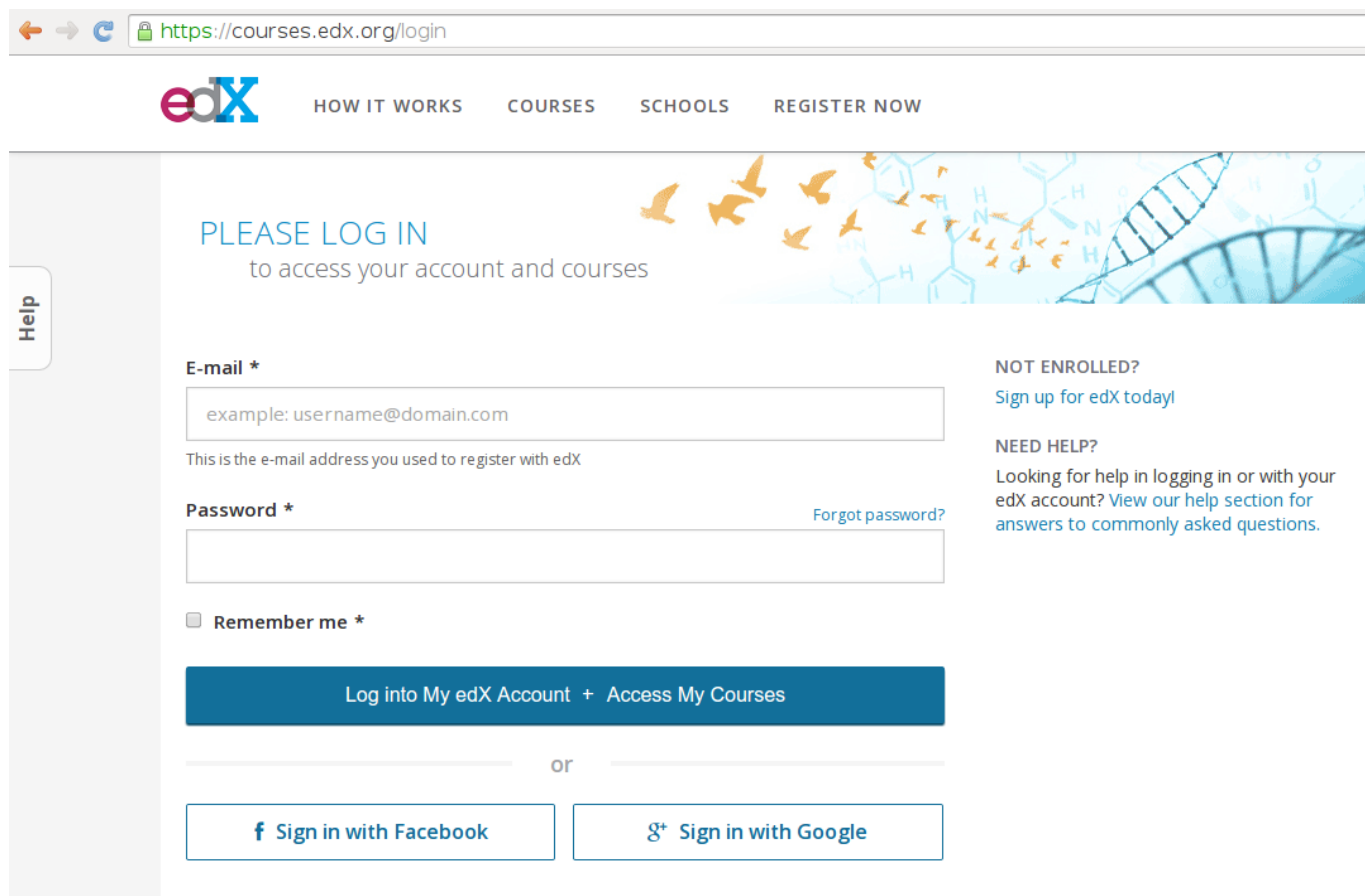
Posted: Thursday, September 18, 2014

Posted by John
Cox, Software
Engineer

 70  46  20

Google is passionate about online education. In addition to our own [Course Builder](#) project, we're also partners with [edX](#), a not-for-profit that shares our desire for scalable, quality education for everyone. Their software, [Open edX](#), lets people make educational content and deliver it online to anybody, anytime, anywhere. It powers their own site, [edx.org](#), and is also used by companies and universities worldwide.

Today we're very pleased to announce that you can now sign in to [edx.org](#) with your Google or Facebook account:

The image is a screenshot of the edX login page in a web browser. The browser's address bar shows the URL 'https://courses.edx.org/login'. The page has a header with the edX logo and navigation links: 'HOW IT WORKS', 'COURSES', 'SCHOOLS', and 'REGISTER NOW'. The main content area is titled 'PLEASE LOG IN to access your account and courses'. It features a 'Help' sidebar on the left. The login form includes an 'E-mail' field with a placeholder 'example: username@domain.com', a 'Password' field, and a 'Remember me' checkbox. A large blue button reads 'Log into My edX Account + Access My Courses'. Below this, there are two buttons for social login: 'f Sign in with Facebook' and 'g+ Sign in with Google'. On the right side of the form, there are links for 'NOT ENROLLED? Sign up for edX today!' and 'NEED HELP? Looking for help in logging in or with your edX account? View our help section for answers to commonly asked questions.' The background of the page features a decorative graphic of orange birds flying over a blue molecular structure.

Until recently, users who wanted to take advantage of the high quality content on edx.org needed to create a new account first. This is a painful, error prone process—really, who wants to worry about yet another password? So we added the ability to use over 60 external authentication providers to Open edX, with support for everything from open standards like [OpenID](#) or [OAuth 2.0](#), to custom university single sign-on systems. For their edx.org site, edX decided to let users pick between Google, Facebook, and a custom username and password.

If you run Open edX, you can also use this feature now. The [authentication module](#) is [extensible](#) so you can add any third-party provider you want if your favorite is not yet supported. And the feature is completely [configurable](#), so you can pick whatever third-party authentication systems are best for your users, including none at all. It's totally up to you.

By simultaneously increasing user choice, convenience, and security, we hope to make open online education even easier and safer to use, whether people pick Course Builder or Open edX for authoring and delivering courses. We're very grateful to our partners at edX for working with us in this exciting field.

Labels: [Education](#), [MOOC](#)

[10 comments](#)

Course Builder now supports the Learning Tools Interoperability (LTI) Specification

Posted: Thursday, September 11, 2014



Posted by John
Cox, Software
Engineer

Since the release of [Course Builder](#) two years ago, it has been used by individuals, companies, and universities worldwide to create and deliver online courses on a variety of subjects, helping to show the potential for making education more accessible through open source technology.

Today, we're excited to announce that Course Builder now supports the [Learning Tools Interoperability](#) (LTI) specification. Course Builder can now interoperate with other LTI-compliant systems and online learning platforms, allowing users to interact with high-quality educational content no matter where it lives. This is an important step toward our goal of making educational content available to everyone.

If you have LTI-compliant software and would like to serve its content inside Course Builder, you can do so by using Course Builder as an LTI consumer. If you want to serve Course Builder content inside another LTI-compliant system, you can use Course Builder as an LTI provider. You can use either of these features, both, or none—the choice is entirely up to you.

The Course Builder LTI extension module, [now available on Github](#), supports LTI version 1.0, and its LTI provider is certified by [IMS Global](#), the nonprofit member organization that created the LTI specification. Like Course Builder itself, this module is open source and available under the Apache 2.0 license.

As part of our continued commitment to online education, we are also happy to announce we have become an affiliate member of IMS Global. IMS Global shares our desire to provide education online at scale, and we look forward to working with the IMS community on LTI and other online education technologies.

Labels: [Education](#), [MOOC](#), [open source](#)

[11 comments](#)

Building a deeper understanding of images

Posted: Friday, September 05, 2014

 947

 422

 1.5k


Posted by

Christian



Szegedy, Software Engineer

The ImageNet large-scale visual recognition challenge ([ILSVRC](#)) is the largest academic challenge in computer vision, held annually to test state-of-the-art technology in image understanding, both in the sense of recognizing objects in images and locating where they are. Participants in the competition include leading academic institutions and industry labs. In 2012 it was won by DNNResearch using the convolutional neural network approach described in the now-seminal [paper by Krizhevsky et al.](#)^[4]

In this year's challenge, team GoogLeNet (named in homage to [LeNet](#), [Yann LeCun](#)'s influential convolutional network) placed first in the classification and detection (with extra training data) tasks, doubling the quality on both tasks over last year's results. The team participated with an open submission, meaning that the exact details of its approach are shared with the wider computer vision community to foster collaboration and accelerate progress in the field.





The competition has three tracks: classification, classification with localization, and detection. The *classification* track measures an algorithm's ability to assign correct labels to an image. The *classification with localization* track is designed to assess how well an algorithm models both the labels of an image and the location of the underlying objects. Finally, the *detection challenge* is similar, but uses much stricter evaluation criteria. As an additional difficulty, this challenge includes a lot of images with tiny objects which are hard to recognize. Superior performance in the detection challenge requires pushing beyond annotating an image with a "bag of labels" -- a model must be able to describe a complex scene by accurately locating and identifying many objects in it. As examples, the images in this post are actual top-scoring inferences of the GoogleNet detection model on the validation set of the detection challenge.



This work was a concerted effort by [Wei Liu](#), [Yangqing Jia](#), [Pierre Sermanet](#), [Scott Reed](#), [Drago Anguelov](#), [Dumitru Erhan](#), [Andrew Rabinovich](#) and [myself](#). Two of the team members -- Wei Liu and Scott Reed -- are PhD students who are a part of the intern program here at Google, and actively participated in the work leading to the submissions. Without their dedication the team could not have won the detection challenge.

This effort was accomplished by using the [DistBelief infrastructure](#), which makes it possible to train neural networks in a distributed manner and rapidly iterate. At the core of the approach is a radically redesigned convolutional network architecture. Its seemingly complex structure (typical incarnations of which consist of over 100 layers with a maximum depth of over 20 parameter layers), is based on two insights: the [Hebbian principle](#) and [scale invariance](#). As the consequence of a careful balancing act, the depth and width of the network are both increased significantly at the cost of a modest growth in evaluation time. The resultant architecture leads to over 10x reduction in the number of parameters compared to most state of the art vision networks. This reduces overfitting during training and allows our system to perform inference with low memory footprint.



For the detection challenge, the improved neural network model is used in the sophisticated [R-CNN detector by Ross Girshick et al.](#)^[2], with additional proposals coming from the [multibox method](#)^[1]. For the classification challenge entry, [several ideas from the work of Andrew Howard](#)^[3] were incorporated and extended, specifically as they relate to image sampling during training and evaluation. The systems were evaluated both stand-alone and as ensembles (averaging the outputs of up to seven models) and their results were submitted as separate entries for transparency and comparison.

These technological advances will enable even better image understanding on our side and the progress is directly transferable to Google products such as photo search, image search, YouTube, self-driving cars, and any place where it is useful to understand *what* is in an image as well as *where* things are.

References:

[1] Erhan D., Szegedy C., Toshev, A., and Anguelov, D., "[Scalable Object Detection using Deep Neural Networks](#)", *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014, pp. 2147-2154.

[2] Girshick, R., Donahue, J., Darrell, T., & Malik, J., "[Rich feature hierarchies for accurate object detection and semantic segmentation](#)", *arXiv preprint arXiv:1311.2524*, 2013.

[3] Howard, A. G., "[Some Improvements on Deep Convolutional Neural Network Based Image Classification](#)", *arXiv preprint arXiv:1312.5402*, 2013.

[4] Krizhevsky, A., Sutskever I., and Hinton, G., "[Imagenet classification with deep convolutional neural networks](#)", *Advances in neural information processing systems*, 2012.

Labels: [Computer Vision](#), [Deep Learning](#), [Neural Networks](#)
[195 comments](#)

Working Together to Support Computer Science Education

Posted: Wednesday, September 03, 2014

Posted by



Tweet

15

Like

15

Chris

Stephenson,

Computer Science Education Program Manager

(Cross-posted from the [Google for Education blog](#))

Computer Science (CS) education in K-12 is receiving an increasing amount of attention from [media](#) and [policy makers](#). Education groups have been working for years to build the infrastructure needed to support CS both inside and outside the school environment, including standards development and dissemination, models for teacher professional development, research, resources for educators, and the building of peer-driven and peer-supported communities of learning.

At Google, we strive to increase opportunities in CS and be a strong contributor to the community of those seeking to improve CS education through our engagement in research, curriculum resource development and dissemination, professional development of teachers, tools development, and large-scale efforts to engage young women and underrepresented groups in computer science. However, despite these efforts, there are still many challenges to overcome to improve the state of CS education.

For example, many people confuse computer science with education technology (the use of computing to support learning in other disciplines) and computer literacy (a very basic understanding of a limited number of computer applications). This confusion leads to the assumption that computer science education is taking place, when in fact in many schools it is not.

Women and minorities are still underrepresented in computer science education and in the high tech workplace. In her introduction to Jane Margolis' [Stuck in the Shallow End: Education, Race, and Computing](#), distinguished scientist Shirley Malcolm refers to computer science as "privileged knowledge" to which minority students often have no access. This statement is supported by data from the [College Board](#) and the [National Center for Women and Information Technology](#).

Poverty also has a significant but often ignored impact on access to technology and quality computer science education. At present there are [more than 16 million U.S. children living in poverty](#); these children are the least likely to have access to computer science knowledge and tools in their schools and homes.

There are many organizations and programs which focus on CS

education, working hard to address these issues, and others. This gives Google the unique opportunity to analyze gaps in existing efforts and apply our resources towards programs that are most needed. In so doing, we hope to help uncover new strategies and create sustainable improvements to CS education.

Achieving systemic and sustained change in K-12 CS education is a complex undertaking that requires strategic support that complements both existing formal school programs and extracurricular education. Google is proud to be a member of the community committed to making tangible improvements to the state of CS education. In future blog posts, we will introduce you so some of the programs and resources that Google has been working on.

Labels: [Computer Science](#), [Education](#), [K-12](#)

[13 comments](#)

Hardware Initiative at Quantum Artificial Intelligence Lab

Posted: Tuesday, September 02, 2014

 94

 80

 253

Posted by
Hartmut Neven,
Director of Engineering

The Quantum Artificial Intelligence team at Google is launching a hardware initiative to design and build new quantum information processors based on superconducting electronics. We are pleased to announce that [John Martinis and his team at UC Santa Barbara](#) will join Google in this initiative. John and his group have made great strides in building [superconducting quantum electronic components of very high fidelity](#). He recently was awarded the [London Prize](#) recognizing him for his pioneering advances in quantum control and quantum information processing. With an integrated hardware group the Quantum AI team will now be able to implement and test new designs for quantum optimization and inference processors based on recent theoretical insights as well as our learnings from the [D-Wave quantum annealing](#) architecture. We will continue to collaborate with D-Wave scientists and to experiment with the “Vesuvius” machine at NASA Ames which will be upgraded to a 1000 qubit “Washington” processor.

Labels: [Machine Learning](#), [Quantum Computing](#)
[20 comments](#)

[Newer Posts](#)

[Home](#)

[Older Posts](#)



- | | | |
|--------------------------------------|-------------------------------|---|
| Company-wide | Products | Developers |
| Official Google Blog | Android Blog | Developers Blog |
| Public Policy Blog | Chrome Blog | Ads Developer Blog |
| Student Blog | Lat Long Blog | Android Developers Blog |